arXiv:1609.06363v1 [math.NA] 20 Sep 2016

EFFICIENT STATIONARY DISTRIBUTION ESTIMATION FOR CONTINUOUS TIME MARKOV CHAIN VIA PARALLEL REPLICA DYNAMICS

TING WANG*, PETR PLECHÁČ†, AND DAVID ARISTOFF‡

Abstract. We propose two algorithms for simulating continuous time Markov chains (CTMC) in the presence of metastability. Both algorithms, based on the idea of the parallel replica method, use parallel computing in order to explore metastable sets more efficiently. The algorithms require no assumptions on the CTMC beyond the presence of identifiable metastability. In particular, there is no assumption on reversibility. We present error analyses, as well as numerical simulations on a simple stochastic reaction network model, to demonstrate consistency of the method and its efficiency.

Key words. Markov chains, Monte Carlo, reversibility, stationary distribution, metastability, long time dynamics, parallel replica, stochastic reaction networks, multiscale dynamics, coarse graining

AMS subject classifications. 60J22, 65C05, 65Z05, 82B31, 92E20

1. Introduction.

We consider the problem of Monte Carlo simulation of continuous time Markov chains in the presence of metastability. Given a stochastic process, a metastable set is a region in state space where the process tends to spend a very long time before leaving. Direct simulation of such processes can be impractical, since simulations can get stuck for long times in the metastable sets.

In computational materials science, Markov processes are often used to model the motion of particles in physical systems. Typically, such models are used to understand a system's macroscopic properties, starting from a microscopic description. Metastability is associated with time scale separation between microscopic and macroscopic events; see [16] for a review of metastability in this setting. The models can be discrete or continuous in time. In this contribution we consider only models described by continuous time Markov chains. As an example, we consider a chemical reaction network model, in which molecules of different types react with different rates depending on their concentrations and reaction rate constants. In this type of model, metastability can arise from reactions with small rates or reactions that become likely only after a buildup of molecules of certain types. Conceptually, metastability arises from entropic barriers, i.e., regions in the state space which can be crossed only through a narrow passageway, as well as energetic barriers, i.e., regions through which crossings are possible but unlikely.

We focus on computing stationary averages of CTMCs. More precisely, if π is the stationary distribution of a CTMC and f is a function on the state space, we aim to compute $\pi(f) \equiv \mathbb{E}_\pi[f]$ by taking a time average of the CTMC. There are many methods for computing stationary averages of stochastic processes, but the vast majority of them rely on reversibility of the process, as in Markov chain Monte Carlo [15]. One advantage of our algorithms is that they may be used, in principle, on arbitrary CTMCs. However, gains in efficiency can occur only if the CTMC is metastable.

Our algorithms use parallel processors to explore metastable sets more efficiently than with direct serial simulation. A similar algorithm for discrete time Markov

*University of Delaware, Newark, DE, 19716 (tingw@udel.edu),
†University of Delaware, Newark, DE, 19716 (plechac@math.udel.edu),
‡Colorado State University, Fort Collins, CO, 80523 (aristoff@rams.colorado.edu).

chains [1, 2] has recently been proposed. See also [12, 18] for a related algorithm for simulating stochastic Langevin dynamics and [13] for a general review. In our algorithms, each time the simulation reaches a local equilibrium in a metastable set W, R replicas of the CTMC are launched inside the set. The main contribution of this work is a procedure for using the replicas in order to *efficiently* and *consistently* estimate the exit time and exit state from W, along with the contribution to the *stationary time average* of f from the time spent in W. We emphasize that we are able to handle arbitrary functions (or observables) on the state space, not just ones that take only one value in each W. In the best case, if there are R replicas, then the simulation leaves a metastable set about R times faster than direct serial simulation. The consistency of our algorithms relies on certain properties of the quai-stationary distributions (QSD) which are essentially local equilibria associated with the metastable sets.

We propose two algorithms for computing $\pi(f)$, called *CTMC ParRep* and *embedded ParRep*. The former uses parallel simulation of the CTMC, while the latter employs parallel simulation of its embedded chain, which is a discrete time Markov chain (DTMC). CTMC ParRep (resp. embedded ParRep) relies on the fact that, starting at the QSD in a metastable set, the first time to leave the set is exponential (resp. geometric) and independent of the exit state; see Theorem 5 below. The algorithms require some methods for identifying metastable sets, though this need not be done a priori – it is sufficient to identify when the CTMC is currently in a metastable set, and when it exits such set. While both algorithms can be useful for efficient simulation of $\pi(f)$ in the presence of metastability, we expect the embedded ParRep can be significantly more efficient, especially when combined with a certain type of QSD sampling, sometimes called Fleming-Viot [3, 4]. Though we focus here on the computation of $\pi(f)$, we note that one of our algorithms, CTMC ParRep, can be used to compute the dynamics of the CTMC on a coarse space in which each metastable set is considered a single (meta-)state. See the discussion below Algorithm 1.

This article is organized as follows. In Section 2, we briefly review CTMCs before defining QSDs and detailing relevant properties thereof. In Section 3, we present CTMC ParRep, and study how the error in the algorithm depends on the quality of QSD sampling. In Section 4, we present embedded ParRep and provide an analogous error analysis. We detail some numerical experiments on a simple chemical reaction network model in Section 5.

2. Preliminaries and Basic Settings.

2.1. Continuous Time Markov Chains. Throughout, $X(t)$ is an irreducible and positive recurrent CTMC with values in a countable set E. Let π be the stationary distribution of $X(t)$. We are interested in computing stationary averages $\pi(f)$ for a bounded function $f : E \to \mathbb{R}$ by using the ergodic theorem

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t f(X(s))ds = \pi(f), \tag{1}$$

which holds almost surely for any initial distribution of $X(t)$. Consider the jump times τ_n and holding times $\Delta\tau_n$ for $X(t)$ defined recursively by

$$\tau_0 = 0, \qquad \tau_n = \inf\{t > \tau_{n-1} : X(t) \neq X(\tau_{n-1})\}, \qquad \Delta\tau_{n-1} = \tau_n - \tau_{n-1} \quad \text{for } n \geq 1.$$

We assume $X(t)$ is non-explosive, that is, $\sum_n \Delta\tau_n = \infty$ almost surely for every initial distribution of $X(t)$. This precludes the possibility of infinitely many jumps in finite

time. We call $X_n = X(\tau_n)$ the embedded chain of $X(t)$. It is easy to see that X_n is a DTMC.

Recall $X(t)$ is completely determined by its infinitesimal generator matrix $Q = \{q(x,y)\}_{x,y\in E}$. We write $q(x) := -q(x,x)$; note that irreducibility implies $q(x) > 0$ for all $x \in E$. It is easy to check that X_n has transition matrix $P = \{p(x,y)\}_{x,y\in E}$ satisfying

$$p(x,y) = \begin{cases} \frac{q(x,y)}{q(x)}, & x \neq y, \\ 0, & x = y \end{cases}.$$

We state the following well known fact for later reference.

LEMMA 1. *For a CTMC $X(t)$ with corresponding embedded Markov chain X_n, the holding time between successive jumps $\Delta\tau_0, \Delta\tau_1, \cdots, \Delta\tau_i, \cdots$ are independent conditioned on the embedded chain X_n. Moreover, $\Delta\tau_i|\{X_n\}$ is exponentially distributed with rate $q(X_i)$ and hence $\mathbb{E}\left[\Delta\tau_i|\{X_n\}\right] = q(X_i)^{-1}$.*

For details on the above facts, see for instance [5].

2.2. The Quasi-stationary Distribution and Metastability. Below, we write $\mathbb{P}$, $\mathbb{E}$ for various probabilities and expectations, the precise meaning of which will be clear from context. We use a superscript $\mathbb{P}^\xi$, $\mathbb{E}^\xi$ to indicate the initial distribution is ξ. When the initial distribution is δ_x, we write $\mathbb{P}^x$, $\mathbb{E}^x$. The symbol $\sim$ will indicate equality in probability law. $\mathrm{Re}(\cdot)$ and $|\cdot|$ denote the real part and modulus of a complex number.

Our ParRep algorithms rely on certain properties of quasi-stationary distributions (QSDs), which we now review. Let $W \subset E$ be fixed. Consider the first exit time of $X(t)$ from W, that is,

$$T = \inf\{t > 0; X(t) \notin W\}.$$

Consider also the first exit time of X_n from W,

$$N = \inf\{n > 0; X_n \notin W\}.$$

A QSD of $X(t)$ in W (or X_n in W) is defined as follows.

DEFINITION 2. *A probability distribution ν with support in W is a QSD for $X(t)$ in W if for each $y \in W$ and $t > 0$,*

$$\nu(y) = \mathbb{P}^\nu(X(t) = y \,|\, T > t). \tag{2}$$

A probability distribution μ with support in W is a QSD for X_n in W if for each $y \in W$ and $n > 0$,

$$\mu(y) = \mathbb{P}^\mu(X_n = y \,|\, N > n). \tag{3}$$

Throughout we write ν for a QSD of the $X(t)$ and μ for a QSD of X_n. The associated set W will be implicit since no ambiguities should arise. We will write

$$\nu_t(A) = \mathbb{P}^x(X(t) \in A \,|\, T > t) \tag{4}$$

for the distribution of $X(t)$ conditioned on $T > t$, and

$$\mu_t(A) = \mathbb{P}^x(X_n \in A \,|\, N > n). \tag{5}$$

for the distribution of X_n conditioned on $N > n$. Notice we do not make explicit the dependence on the starting point x.

We summarize existence, uniqueness, and convergence properties of the QSD in Theorem 3 below (see [6, 17]). In Theorem 3 below, for simpler presentation, we assume W is finite. That allows us to characterize convergence to the QSD of $X(t)$ and X_n in terms of spectral properties their generator and transition matrices. We emphasize, however, that finiteness of W is not required for consistency of the algorithms proposed in this paper.

Let $Q_W = \{q_{xy}\}_{x,y\in W}$ and $P_W = \{p_{xy}\}_{x,y\in W}$ be the restrictions of P and Q to W.

THEOREM 3. *Let W be finite and nonabsorbing for $X(t)$, and assume P_W is irreducible.*

(a) The eigenvalues $\lambda_1, \lambda_2, \ldots$ of Q_W can be ordered so that

$$0 < \lambda_1 < \mathrm{Re}(\lambda_2) \leq \ldots,$$

where λ_1 has left eigenvector ν which is a probability distribution on W. Moreover, ν is the unique QSD for $X(t)$ in W, and for all $x, y \in W$,

$$|\nu_t(y) - \nu(y)| = |\mathbb{P}^x(X(t) = y | T > t) - \nu(y)| \leq C(x) e^{-(\lambda_1 - \beta)t}, \tag{6}$$

with $C(x)$ a constant depending on x, and β any real number satisfying $\mathrm{Re}(\lambda_2) < \beta < \lambda_1$.

(b) Suppose P_W is also aperiodic. Then the eigenvalues $\sigma_1, \sigma_2, \ldots$ of P can be ordered so that

$$1 > \sigma_1 > |\sigma_2| \geq \ldots,$$

where σ_1 has left eigenvalue μ which is a probability distribution on W. Moreover, μ is the unique QSD of X_n in W and for all $x, y \in W$,

$$|\mu_t(y) - \mu(y)| = |\mathbb{P}^x(X_n = y | N > n) - \mu(y)| \leq D(x) \left(\frac{\gamma}{\sigma_1}\right)^n, \tag{7}$$

with $D(x)$ a constant depending on x, and γ any real number satisfying $\gamma > |\sigma_2|$.

Proof. We first justify the expression for the eigenvalues. Observe that for $x \neq y \in W$, we have $q(x, y) > 0$ if and only if $p(x, y) > 0$. It follows that Q_W is irreducible if and only if P_W is irreducible; see Definition 2.1 in [17]. Now let e be the all ones column vector, $e(x) = 1$ for $x \in W$. Recall that $q(x, y) \geq 0$ for every $x \neq y \in E$ and $\sum_y q(x, y) = 0$ for every $x \in E$. This implies that $Q_W e \leq 0$ componentwise. Since W is nonabsorbing, for some $x \in W$ and $y \notin W$ we have $q(x, y) > 0$, and it follows that $\sum_{z\in W} q(x, z) < 0$. This shows that at least one component of $Q_W e$ is strictly negative. The expression for the eigenvalues, and the fact that ν is signed (hence a probability distribution, after normalization) now follows from Theorem 2.6 of Seneta [17].

To see ν is the QSD for $X(t)$ in W, define the stopped process $X^T(t) = X(t \wedge T)$ such that $X(t)$ is absorbed outside W. For $x, z \in E$, let e_x be the column vector $e_x(z) = 1$ if $x = z$ and $e_x(z) = 0$ otherwise. Finiteness of W ensures that $\mathbb{P}^x(X^T(t) = y) = e_x^T e^{Q_W t} e_y$. Thus, for each $y \in W$,

$$\mathbb{P}^\nu(X(t) = y, T > t) = \mathbb{P}^\nu(X^T(t) = y) = \nu e^{Q_W t} e_y = e^{\lambda_1 t} \nu(y)$$

and

$$\mathbb{P}^\nu(T > t) = \mathbb{P}^\nu(X^T(t) \in W) = e^{\lambda_1 t},$$

which leads to $\nu(y) = \mathbb{P}^\nu(X(t) = y | T > t)$.

Now we turn to the convergence to ν. It follows from Theorem 2.7 in [17] that there is a constant $C(x)$ depending on x such that for any real β with $\mathrm{Re}(\lambda_2) < \beta$,

$$\mathbb{P}^x(X(t) = y, T > t) = \mathbb{P}^x(X^T(t) = y) = C(x)e^{\lambda_1 t}\nu(y) + \mathcal{O}(e^{\beta t}) \tag{8}$$

and

$$\mathbb{P}^x(T > t) = C(x)e^{\lambda_1 t} + \mathcal{O}(e^{\beta t}), \tag{9}$$

It follows that

$$|\nu_t(t) - \nu(y)| = |\mathbb{P}^x(X(t) = y \,|\, T > t) - \nu(y)| \le C(x)e^{-(\lambda_1 - \beta)t}$$ ∎

where $C(x)$ is now a (possibly different) constant depending on x.

The arguments in (b) are similar, using the Perron-Frobenius theorem (Theorem 1.1 in Seneta [17]).

For analogous results on the QSD in more general settings, see [6] (Theorem 4.5) for CTMCs and to [8] (Theorem 1) for DTMCs. We are now ready to define metastability.

Definition 4. *Let W and λ_i, σ_i be as in Theorem 3.*

1. *W is metastable for $X(t)$ if $\lambda_1 \approx 0$ and*

$$|\lambda_1| \ll |\lambda_1 - \mathrm{Re}(\lambda_2)|. \tag{10}$$

   $X(t)$ is metastable if it has at least one metastable set W.

2. *W is metastable for X_n if $\sigma_1 \approx 1$ and*

$$1 - \sigma_1 \ll 1 - \frac{|\sigma_2|}{\sigma_1}. \tag{11}$$

   X_n is metastable if it has at least one metastable set W.

In light of Theorem 3, conditions 1-2 in Definition 4 essentially say that the time to leave W is large in an absolute sense, and the time to leave W is large relative to the time to converge to the QSD in W. Metastability of the CTMC is not necessarily equivalent to the metastability of its underlying embedded chain, as we now show. Consider $X(t)$ with infinitesimal generator

$$Q = \begin{pmatrix} -1 & 1/2 & 1/2 & 0 \\ 1/2 & -1 & 1/2 & 0 \\ 0 & \epsilon/2 & -\epsilon & \epsilon/2 \\ 0 & 0 & 1 & -1 \end{pmatrix},$$

where $\epsilon \approx 0$ is positive. Then $W = \{1, 2, 3\}$ is metastable for $X(t)$ but not for X_n, since

$$\sigma_1 \approx 0.81, \quad |\sigma_2| \approx 1/2, \qquad \lambda_1 \approx -\epsilon/2, \quad \mathrm{Re}(\lambda_2) \approx -1/2.$$

Now consider $X(t)$ with infinitesimal generator

$$Q = \begin{pmatrix} -\epsilon^{-1} & \epsilon^{-1}/2 & \epsilon^{-1}/2 & 0 \\ \epsilon^{-1} - 1 & -\epsilon^{-1} & 1 & 0 \\ 0 & \epsilon^{-1} - 1 & -\epsilon^{-1} & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Then $W = \{1, 2, 3\}$ is metastable for X_n but not for $X(t)$, since

$$\sigma_1 \approx 1 - \epsilon/5, \quad \sigma_2 \approx -1/2 \pm i/2, \qquad \lambda_1 \approx -1/5, \quad \lambda_2 \approx -3\epsilon^{-1}/2 \pm i\epsilon^{-1}/2.$$

Algorithm 1 below requires a collection of metastable sets for $X(t)$, and Algorithm 2 requires a collection of metastable sets for X_n. The only assumption we make on these sets is that they are pairwise disjoint. (The sets may be different for the two algorithms, as noted above.) Throughout write W to denote a generic metastable set. We emphasize that *we do not assume the metastable sets form a partition of* E: the union of the metastable sets may be a proper subset of E. Here and below, we assume that each W has a unique QSD and that ν_t (and μ_t) converge to the QSD in total variation norm, for any starting point x.

We conclude this section by mentioning properties of the QSD which are essential for the consistency of our algorithms in Section 3 and 4 below.

THEOREM 5.

1. *Suppose* $X(0) \sim \nu$. *Then* T *is exponentially distributed with parameter* λ_1*:*

$$\mathbb{P}^\nu(T > t) = e^{\lambda_1 t}, \qquad t > 0,$$

and T *and* $X(T)$ *are independent.*

2. *Suppose* $X_0 \sim \mu$. *Then* N *is geometrically distributed with parameter* $1 - \sigma_1$*:*

$$\mathbb{P}^\mu(N > n) = \sigma_1^n, \qquad n = 1, 2, \ldots,$$

and N *and* X_N *are independent.*

Proof. The first part of 1 and 2 was shown in Theorem 3. For the rest of the proof see [6]. ∎

3. The CTMC ParRep Method.

3.1. Formulation of the CTMC Algorithm . In this section, we introduce a method for accelerating the computation of $\pi(f)$, where we recall $f : E \to \mathbb{R}$ is any bounded function. We call this algorithm CTMC ParRep, for reasons that will be outlined below. Before we describe CTMC ParRep, we introduce some notation. Throughout, $X^1(t), \ldots, X^R(t)$ will be independent processes with the same law as $X(t)$ and with initial distributions supported in W. Recall that the first exit time of $X(t)$ from W is

$$T = \inf\{t > 0 : X(t) \notin W\}.$$

For $r = 1, \ldots, R$, we define the first exit time of $X^r(t)$ from W by

$$T^r = \inf\{t > 0 : X^r(t) \notin W\}$$

and the smallest one among them by

$$T^* = \min_r T^r.$$

We denote the index of the replica with first exit time T^* by M, i.e.,

$$M = \arg\min_r T^r.$$

Of course T, T^r, T^* and M depend on W, but we do not make this explicit.

Algorithm 1 CTMC ParRep

1: Set a decorrelation threshold t_c for each metastable set W. Initialize the simulation time clock $T_{\rm sim} = 0$ and the accumulated value $F(f)_{\rm sim} = 0$. We will write $X^{\rm par}(t)$ for a simulation process that obeys the law of $X(t)$. A complete ParRep cycle consists of three stages.

2: **Decorrelation Stage** : Starting at $t = T_{sim}$, evolve $X^{\rm par}(t)$ until it spends an interval of time length t_c inside of the same metastable set W. That is, evolve $X^{\rm par}(t)$ from time $t = T_{sim}$ until time

$$T_{corr} = \inf\{t \geq T_{sim} + t_c : \exists\ W\ s.t. X^{par}(s) \in W\, \forall\, s \in [t - t_c, t]\}.$$

Then update

$$F(f)_{sim} = F(f)_{sim} + \int_{T_{sim}}^{T_{corr}} f(X^{\rm par}(t))\, dt,$$

set $T_{sim} = T_{corr}$, and go to the dephasing stage.

3: **Dephasing Stage** : Let W be such that $X^{\rm par}(T_{sim}) \in W$, that is, W is the metastable set from the end of the decorrelation stage. Generate R independent samples $x_1, \ldots, x_R$ from ν, the QSD of $X(t)$ in W. Then go to the parallel stage.

4: **Parallel Stage** : Start R parallel processes $X^1(t), \ldots, X^R(t)$ at $x_1, \ldots, x_R$, and evolve them from time $t = 0$ until time T^*. Then update

$$\begin{aligned} F(f)_{\rm sim} &= F(f)_{\rm sim} + \sum_{r=1}^{R} \int_0^{T^*} f(X^r(s))ds, \\ T_{sim} &= T_{sim} + RT^*, \end{aligned} \tag{12}$$

set $X^{\rm par}(T_{sim}) = X^M(T^*)$, and return to the decorrelation stage.

5: The algorithm is stopped when $T_{\rm sim}$ reached some user-chosen time $T_{\rm end}$. The stationary average $\pi(f)$ is estimated as

$$\pi(f) \approx F(f)_{\rm sim}/T_{\rm sim}.$$

We are in position to present Algorithm 1 - CTMC ParRep. In this algorithm, we will need user-chosen parameters t_c associated with each metastable set W. Roughly, these parameters correspond to the time for $X(t)$ to converge to the QSD in W. A few comments are in order.

- The CTMC $X(t)$ can be simulated by the classical stochastic SSA (see [10] for reference). Essentially, one samples the holding times $\Delta\tau_n$ and the next position of the embedded chain X_{n+1} at the same time, using two random numbers. The process $X(t)$ can then be reconstructed using

$$X(t) = X_n, \qquad \Delta\tau_0 + \ldots + \Delta\tau_n \leq t < \Delta\tau_0 + \ldots + \Delta\tau_{n+1}.$$

  This is how we simulate $X^{\rm par}(t)$ and $X^1(t), \ldots, X^R(t)$ in the decorrelation and parallel stages.
- If $X^{\rm par}(t)$ remains in W for sufficiently long time (i.e., time t_c), it is dis-

tributed nearly according to the QSD ν of $X(t)$ in W. See Theorem 3. This means that at the end of the decorrelation stage, $X^{\text{par}}(T_{sim})$ can be considered a sample of ν.

- The aim of the dephasing stage is to prepare a sequence of i.i.d. initial states with distribution ν. There are many possibilities for doing this. Perhaps the simplest is the rejection method. In this procedure, each of the R replicas evolves independently. A parameter t_p similar to the decorrelation time t_c is selected. If a replica leaves W before spending a time interval of length t_p in W, it restarts in W. Once all the replicas remain in W for time t_p, we stop and take $x_1, \dots, x_R$ as the final states of all the replicas. Besides rejection sampling, another possibility is a Fleming-Viot based particle sampler; see the discussion after Algorithm 2 below.
- The acceleration of CTMC ParRep comes from the parallel stage. Recall that, for each $r = 1, \dots, R$, if $x_1, \dots, x_R$ are iid with common distribution ν, then $T^1, \dots, T_R$ are independent exponential random variables with common parameter λ_1. Using $T^* = \min_r T^r$, it is then easy to check that RT^* has the same distribution as T^1. See Lemma 6 below.
  This means one only needs to wait for T^* instead of T^1 to observe an exit from W. Note that this is true whether or not W is metastable, so efficiency of the parallel stage does not require metastability. However, the dephasing step is not efficient if W is not metastable. That is because, in practice, the samples $x_1, \dots, x_R$ are obtained by simulating trajectories which remain in W for a sufficiently long time t_p. Such samples are hard to obtain when the typical time t_p for $x_1, \dots, x_R$ to reach the QSD in W is not much smaller than the typical time to leave W.
- To see that each parallel stage has a consistent contribution to the stationary average, we make the following two observations. Suppose that $x_1, \dots, x_R$ are iid samples from ν. Then:
  1. The joint law of $(RT^*, X^M(T^*))$ is the same as that of $(T^1, X(T^1))$. That is, the joint distribution of the first exit time and the exit state in the parallel stage is independent of the number of replicas.
  2. The expected value of $\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s))ds$ is the same as that of $\int_0^{T^1} f(X^1(s))ds$. That is, the expected contribution to $F(f)_{\text{sim}}$ from each parallel stage is independent of the number of replicas.

  The first observation a consequence of the Theorem 5, and the second will be proved in Theorem 7 below. Consistency of stationary averages follows from 1-2 above and the law of large numbers: since there are indefinitely many parallel stages in a given W, consistency is ensured so long as the expected contribution to $F(f)_{sim}$ from the parallel stage has the correct expected value. (See [1] for details and discussion in a related discrete time version of the algorithm.) It isn't necessary that the contribution to $F(f)_{sim}$ from the parallel stage has the same *law* for any number of replicas – though this turns out to be true for CTMC ParRep.

The CTMC ParRep algorithm suffers some serious drawbacks. Even if the parallel processors are synchronous, M and T^* may not be known at the wall clock time when the first replica leaves W. The reason is that the holding times for a CTMC are random, while the wall clock time for simulating each jump of the CTMC is always roughly the same. We illustrate this problem in Figure 1. In the worst possible case, in order to determine M and T^*, we must wait for all the replicas to leave

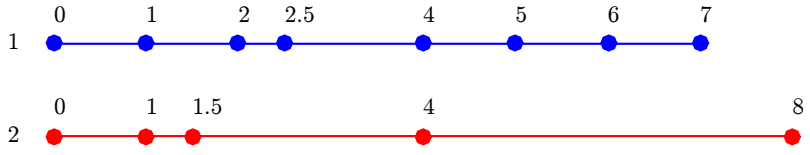

FIG. 1. *The parallel stage of the CTMC ParRep algorithm with two replicas. $R1$ escapes from W at $t = 7$ with 7 transitions while $R2$ escapes at $t = 8$ but with only 4 transitions. In the parallel stage of the CTMC ParRep algorithm, $R2$ escaped from W before $R1$ does but $T^2 > T^1$. There is no acceleration in this case since the parallel stage does not terminate when $R2$ escapes.*

W. However, one can set a variable $T_{\min}$ to record the current minimum first exit time over all replicas which have left W, and terminate any replicas which reach time $T_{\min}$ but have not left W, since no replica contributes to the accumulated value past time $T_{\min}$. Since the expected first exit times $\mathbb{E}[T^r], r = 1, \ldots, R$ are roughly the same, if the variance in the number of jumps of $X^r(t)$ before time T^* is small for all $r = 1, \ldots, R$, then we can expect that the parallel stage stops after only a few replicas leave W.

There is another major drawback of CTMC ParRep. If f takes multiple values in W, then the computation of $\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s))ds$ in (12) requires the storing entire history of each replica in that parallel stage. Hence, the implementation of the CTMC ParRep might be memory intense unless one is interested in the equilibrium average of a metastable-set invariant function f, i.e., if $f(x)$ has only one value inside each metastable set W. In Section 4 we present another algorithm, called embedded ParRep, which addresses these drawbacks.

3.2. Error Analysis of CTMC ParRep. Here and below we will write $\mathbb{E}^{\nu^R}$ for the expectation of $(X^1(t), \ldots, X^R(t))$ starting at ν^R, where for

$$\nu^R(x_1, \ldots, x_R) = \prod_{r=1}^{R} \nu(x_r), \qquad x_1, \ldots, x_R \in W.$$

We begin with a simple well known lemma.

LEMMA 6. *Suppose $T^1, \ldots, T^R$ are iid exponential random variables with parameter λ_1. Then $T^* := \min_{1 \le r \le R} T^r$ is exponentially distributed with parameter $R\lambda_1$. In particular, RT^* has the same distribution as T^1.*

We now show that if the dephasing sampling is exact, then the expected contribution to the accumulated value $F(f)_{sim}$ from the parallel step of Algorithm 1 is exact.

THEOREM 7. *Suppose in the dephasing step $(x_1, \ldots, x_R) \sim \nu^R$. Then the expected contribution to $F(f)_{sim}$ from the parallel stage of Algorithm 1 is the same for every number of replicas:*

$$\mathbb{E}^{\nu^R}\left[\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s))ds\right] = \mathbb{E}^{\nu}\left[\int_0^{T} f(X(s))ds\right] = \nu(f)\mathbb{E}^{\nu}[T].$$

Proof. We condition on the exit time T^1 and write

$$\mathbb{E}^{\nu}\left[\int_0^{T^1} f(X^1(s))ds\right] = \int_0^{\infty} \mathbb{E}^{\nu}\left[\int_0^{t} f(X^1(s))ds \,\middle|\, T^1 = t\right] \mathbb{P}^{\nu}(T^1 \in dt).$$

Interchanging the two integrals of the right-hand side leads to

$$\int_0^\infty \int_s^\infty \mathbb{E}^\nu \left[f(X^1(s)) | T^1 = t \right] \mathbb{P}(T^1 \in dt) ds.$$

Note that the inner integral can be written as $\mathbb{E}^\nu \left[f(X^1(s)) \mathbb{1}_{T^1 > s} \right]$ and hence

$$\mathbb{E}^\nu \left[\int_0^{T^1} f(X^1(s)) ds \right] = \int_0^\infty \mathbb{E}^\nu \left[f(X^1(s)) | T^1 > s \right] \mathbb{P}^\nu (T^1 > s) ds.$$

Owing to the definition of QSD and the fact that $\mathbb{E}^\nu[T^1] = \int_0^\infty \mathbb{P}^\nu(T^1 > s) ds$, hence

$$\mathbb{E}^\nu \left[\int_0^{T^1} f(X^1(s)) ds \right] = \nu(f) \mathbb{E}^\nu[T^1].$$

In the case of multiple replicas, similar steps can be used to show that

$$\sum_{r=1}^R \mathbb{E}^{\nu^R} \left[\int_0^{T^*} f(X^r(s)) ds \right] = \sum_{r=1}^R \int_0^\infty \mathbb{E}^{\nu^R} \left[f(X^r(s)) | T^* > s \right] \mathbb{P}^{\nu^R} (T^* > s) ds.$$

Recall that $T^* > s$ if and only if $T^r > s$ for all $r = 1, \ldots, R$. Using this, the fact that $T^1, \ldots, T^r$ are independent, and the definition of the QSD, we get

$$\mathbb{E}^{\nu^R} \left[f(X^r(s)) | T^* > s \right] = \mathbb{E}^\nu \left[f(X^r(s)) | T^r > s \right] = \nu(f).$$

Thus

$$\sum_{r=1}^R \mathbb{E}^{\nu^R} \left[\int_0^{T^*} f(X^r(s)) ds \right] = \nu(f) \sum_{r=1}^R \int_0^\infty \mathbb{P}^{\nu^R} (T^* > s) ds = \nu(f) R \mathbb{E}^{\nu^R}[T^*].$$

Finally by the last equality coming from Lemma 6,

$$\nu(f) R \mathbb{E}^{\nu^R}[T^*] = \nu(f) \mathbb{E}^\nu[T]. \qquad \square$$

The purpose of CTMC ParRep is to efficiently simulate very long trajectories of a metastable CTMC and estimate the equilibrium average $\pi(f)$. (CTMC ParRep can produce accelerated dynamics of the CTMC on a coarse state space where each coarse set corresponds to some W; see the discussion below Algorithm 2 below.) Our numerical experiments suggest that CTMC ParRep (and also embedded ParRep described below) are consistent for estimating the stationary distribution. For CTMC ParRep, we justify this claim in Theorem 8 below, which shows that, starting in some W and waiting until the simulation leaves W, the error in CTMC ParRep compared to direct (serial) simulation vanishes as t_c increases. See Theorem 12 below for the analogous result on embedded ParRep. We note that the errors from each ParRep cycle produce an error in the estimation (5) of stationary averages that does not disappear as $T_{sim} \to \infty$. Study of the latter error is more involved and will be the focus of another work.

Recall we have assumed convergence of $\|\nu_{t_c} - \nu\|_{TV} \to 0$ as $t_c \to \infty$, for every starting point $x \in E$, where $\| \cdot \|$ denotes total variation norm. See for instance Theorem 3 for conditions guaranteeing this convergence.

THEOREM 8. *Consider CTMC ParRep starting at $X^{par}(T_{sim}) = x \in W$ in the decorrelation stage, where WLOG $T_{sim} = 0$. Define*

$$T^{par} = \begin{cases} T & \text{if } t_c > T \\ t_c + RT^* & \text{if } t_c \le T. \end{cases} \tag{13}$$

Assume the dephasing stage sampling is exact, that is, $(x_1, \dots, x_R) \sim \nu^R$. Consider the expected contribution to $F(f)_{sim}$ up until the first time the simulation leaves W (either in the decorrelation or in the parallel stage):

$$\Delta F(f)_{sim} \triangleq \mathbb{E}^{x,\nu^R}\left[\int_0^{t_c \wedge T^{par}} f(X^{par}(s))\, ds + \mathbb{1}_{T^{par} > t_c} \sum_{r=1}^{R} \int_0^{T^*} f(X^r(s)) ds\right],$$

where $\mathbb{E}^{x,\nu^R}$ denotes expectation corresponding to $X^{par}(t)$ starting at x and the replicas $(X^1(t), \dots, X^R(t))$ with initial distribution ν^R. The error compared to direct (serial) simulation satisfies the bound

$$\left|\mathbb{E}^x\left[\int_0^T f(X(s)) ds\right] - \Delta F(f)_{sim}\right| \le \|f\|_\infty \sup_{x \in W} \mathbb{E}^x[T]\, \|\nu_{t_c} - \nu\|_{TV}. \tag{14}$$

Before proving the theorem, we note that $\mathbb{E}^x[T]$ is uniformly bounded in $x \in W$ if, for instance, P_W is irreducible and W is finite and nonabsorbing for $X(t)$, as in Theorem 3. This uniform boundedness guarantees that the right hand side of (14) vanishes as $t_c \to \infty$.

Proof of Theorem 8. Since $X^{par}(t)$ has the same law as $X(t)$,

$$\mathbb{E}^x\left[\int_0^{t_c \wedge T} f(X(s)) ds\right] = \mathbb{E}^{x,\nu^R}\left[\int_0^{t_c \wedge T^{\mathrm{par}}} f(X^{\mathrm{par}}(s))\, ds\right].$$

Thus,

$$\begin{aligned}
&\left|\mathbb{E}^x\left[\int_0^T f(X(s)) ds\right] - \Delta F(f)_{sim}\right| \\
&= \left|\mathbb{E}^x\left[\int_{t_c \wedge T}^T f(X(s)) ds\right] - \mathbb{E}^{x,\nu^R}\left[\mathbb{1}_{T^{par} > t_c} \sum_{r=1}^{R} \int_0^{T^*} f(X^r(s)) ds\right]\right| \\
&= \left|\mathbb{E}^x\left[\int_{t_c}^T f(X(s)) ds \,\middle|\, T > t_c\right] - \mathbb{E}^{x,\nu^R}\left[\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s)) ds \,\middle|\, T^{\mathrm{par}} > t_c\right]\right| \mathbb{P}^x(T > t_c) \\
&\le \left|\mathbb{E}^x\left[\int_{t_c}^T f(X(s)) ds \,\middle|\, T > t_c\right] - \mathbb{E}^{\nu^R}\left[\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s)) ds\right]\right|.
\end{aligned}$$

By the Markov property,

$$\mathbb{E}^x\left[\int_{t_c}^T f(X(s)) ds \,\middle|\, T > t_c\right] = \mathbb{E}^{\nu_{t_c}}\left[\int_0^T f(X(s)) ds\right].$$

By Theorem 7,

$$\mathbb{E}^{\nu^R}\left[\sum_{r=1}^{R} \int_0^{T^*} f(X^r(s)) ds\right] = \mathbb{E}^{\nu}\left[\int_0^T f(X(s))\, ds\right].$$

Combining the last three equations,

$$\begin{aligned}
&\left|\mathbb{E}^x\left[\int_0^T f(X(s))ds\right] - \Delta F(f)_{sim}\right| \\
\leq &\left|\mathbb{E}^{\nu_{t_c}}\left[\int_0^T f(X(s))ds\right] - \mathbb{E}^{\nu}\left[\int_0^T f(X(s))\,ds\right]\right| \\
= &\left|\sum_{x\in W}\mathbb{E}^x\left[\int_0^T f(X(s))ds\right]\nu_{t_c}(x) - \sum_{x\in W}\mathbb{E}^x\left[\int_0^T f(X(s))ds\right]\nu(x)\right| \\
\leq &\|f\|_\infty \sup_{x\in W}\mathbb{E}^x\left[T\right]\|\nu_{t_c}-\nu\|_{\mathrm{TV}}.
\end{aligned}$$

∎

4. The Embedded ParRep Method.

4.1. Formulation of the Embedded ParRep Algorithm. In this section, we introduce another algorithm for accelerating the computation of $\pi(f)$. The algorithm, called embedded ParRep, circumvents the disadvantages of CTMC ParRep discussed above. As mentioned in the previous section, CTMC ParRep can be slow due to the randomness of the holding times. In the worst case, one has to wait until all replicas leave W to determine the first exit time T^*. To circumvent this issue, we propose an algorithm based on the embedded chain, in which the parallel stage terminates as long as one of the replicas leaves W.

Before we describe embedded ParRep, we introduce some notations. Throughout, $X_n^1,\ldots,X_n^R$ will be independent processes with the same law as X_n and with initial distributions supported in W. Moreover, we consider $X_n^1,\ldots,X_n^R$ as the embedded chains of $X^1(t),\ldots,X^r(t)$ defined above, and let $\Delta\tau_n^1,\ldots,\Delta\tau_n^R$ be the corresponding holding times. Recall that the first exit time of X_n from W is

$$N = \inf\{n>0 : X_n \notin W\}.$$

For $r=1,\ldots,R$, we define the first exit time of X_n^r from W by

$$N^r = \min\{n\in\mathbb{N}; X_n^r \notin W\}$$

and the smallest among them by

$$N^* = \min\{N^r; r=1,\ldots,R\}.$$

Note that it is possible that more than one replicas leaves W for the first time after N^* transitions. We denote by K the smallest index among these replicas. That is,

$$K = \min\{r=1,\ldots,R; X_{N^*}^k \notin W\}.$$

It is clear from the above definition that $N^K = N^*$. Of course N, N^r, N^* and K depend on W, but we do not make this explicit.

Here and below we write $\mathbb{E}^{\mu^R}$ for expectation of $(X_n^1,\ldots,X_n^R)$ starting at μ^R, where

$$\mu^R(x_1,\ldots,x_R) = \prod_{r=1}^R \mu(x_r), \qquad x_1,\ldots,x_R\in W.$$

We begin by reproducing from [2] Theorems 9 and 10 below, with proofs for completeness.

THEOREM 9. *Suppose* $(X_n^1, \dots, X_n^R)$ *has initial distribution* μ^R. *Then* $R(N^K - 1) + K$ *has the same distribution as* N^1.

Proof. By Theorem 5, N_1 is geometrically distributed with rate $\mathbb{P}^\mu(N^1 > 1)$. Note that for any $n \geq 0$ and $r = 1, \dots, R$, the event $\{N^K = n, K = k\}$ is equivalent to the event $\{N^1 > n, \dots, N^{k-1} > n, N^k = n, N^{k+1} > n-1, \dots, N^R > n-1\}$. Since $X_n^1, \dots, X_n^R$ are iid and N^1 is geometrically distributed with rate $p = \mathbb{P}^\mu(N^1 > 1)$,

$$\mathbb{P}(N^K = n, K = k) = (1-p)^{n(k-1)}(1-p)^{n-1}p(1-p)^{(n-1)(k-1)} = (1-p)^{R(n-1)+k-1}p.$$

That is, $R(N^K - 1) + K$ has geometric distribution with rate p. ∎

THEOREM 10. *Suppose* $(X_n^1, \dots, X_n^R)$ *has initial distribution* μ^R. *Then* $X_{N^K}^K$ *is independent of* $R(N^K - 1) + K$ *and the distribution of* $(X_{N^K}^K, R(N^K - 1) + K)$ *is same as that of* $(X_{N^1}^1, N^1)$.

Proof. We first prove that $X_{N^K}^K$ is independent of K. Since $X_n^R, \dots, X_n^R$ are iid and N^k is independent of $X_{N^k}^k$ for each k, then $X_{N^k}^k$ is independent of $N^1, \dots, N^R$. Note that $K \in \sigma(N^1, \dots, N^R)$, hence $X_{N^k}^k$ is independent of K for each k. Now observe that for any $A \subset E$,

$$\begin{aligned}
\mathbb{P}(X_{N^K}^K \in A) &= \sum_{r=1}^R \mathbb{P}(X_{N^k}^k \in A, K = r) \\
&= \sum_{r=1}^R \mathbb{P}(X_{N^1}^1 \in A)\mathbb{P}(K = r) \\
&= \mathbb{P}(X_{N^1}^1 \in A),
\end{aligned}$$

that is, $X_{N^K}^K$ and $X_{N^1}^1$ are equally distributed. This implies that $X_{N^K}^K$ is independent of K. To see the independence between $X_{N^K}^K$ and $R(N^K - 1) + K$, note that

$$\begin{aligned}
\mathbb{P}(X_{N^K}^K \in A, N^K = n, K = r) &= \mathbb{P}(X_{N^r}^r \in A, N^r = n, K = r) \\
&= \mathbb{P}(X_{N^r}^r \in A, K = r | N^r = n)\mathbb{P}(N^r = n) \\
&= \mathbb{P}(X_{N^r}^r \in A | N^r = n)\mathbb{P}(N^r = n, K = r) \\
&= \mathbb{P}(X_{N^r}^r \in A)\mathbb{P}(N^r = n, K = r) \\
&= \mathbb{P}(X_{N^K}^K \in A)\mathbb{P}(N^K = n, K = r)
\end{aligned}$$

for any measurable $A \subset E$, $n \in \mathbb{Z}^+$ and $r = 1, \dots, R$. Finally, Theorem 9 and the above analysis imply that $(X_{N^K}^K, R(N^K - 1) + K)$ and $(X_{N^1}^1, N^1)$ are equally distributed. ∎

In Algorithm 2 we will need user-chosen parameters n_c associated with each metastable set W. Roughly, these parameters correspond to the time for X_n to converge to the QSD in W.

We make the following comments.

- The DTMC X_n and holding times $\Delta\tau_n$ are simulated by the classical stochastic SSA (see [10] for instance), just as in CTMC ParRep. If X_n^{par} remains in W for sufficiently long time (i.e., time t_c), it is distributed nearly according to the QSD μ of X_n in W. See Theorem 3. This means that at the end of the decorrelation stage, X_n^{par} can be considered a sample of μ.
- The aim of the dephasing stage is to prepare a sequence of i.i.d. initial states with distribution μ. There are many possibilities for doing this. One possibility is rejection sampling. But perhaps the most natural option for

Algorithm 2 Embedded ParRep

1: Set a decorrelation threshold n_c for each metastable set W. Initialize the the simulation time clock $N_{\rm sim} = 0$ and the accumulated value $F(f)_{\rm sim} = 0$. We will write X_n^{par} and Δ_n^{par} for a DTMC/holding time process following the law of the embedded chain/holding times of $X(t)$. A complete ParRep cycle consists of three stages.

2: **Decorrelation Stage**: Starting at $n = N_{sim}$, evolve X_n^{par} and $\Delta\tau_n^{par}$ until X_n^{par} spends n_c consecutive time steps inside of the same metastable set W. That is, evolve X_n^{par} and $\Delta\tau_n^{par}$ from time $n = N_{sim}$ until time

$$N_{corr} = \inf\{n \geq N_{sim} + n_c - 1 : \exists\, W \, s.t. \, X_m^{par} \in W \, \forall\, m \in \{n - n_c + 1, \ldots, n\}\}.$$

Then update

$$F(f)_{sim} = F(f)_{sim} + \sum_{n=N_{sim}}^{N_{corr}-1} f(X_n^{par})\Delta\tau_n^{par},$$

set $N_{sim} = N_{corr}$, and go to the dephasing stage.

3: **Dephasing Stage** : Let W be such that $X_{N_{sim}}^{par} \in W$, that is, W is the metastable set from the end of the decorrelation stage. Generate R independent samples $x_1, \ldots, x_R$ from μ, the QSD of X_n in W. Then proceed to the parallel stage.

4: **Parallel Stage** : Start R parallel processes $X_n^1, \ldots, X_n^R$ at $x_1, \ldots, x_R$, and evolve them and the corresponding holding times $\Delta\tau_n^1, \ldots, \Delta\tau_n^R$ from time $n = 0$ until time N^*. Then update

$$\begin{aligned} F(f)_{\rm sim} &= F(f)_{\rm sim} + \sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r \\ N_{sim} &= N_{sim} + R(N^* - 1) + K, \end{aligned} \tag{15}$$

set $X_{N_{sim}}^{\rm par} = X_{N^*}^K$, and return to the decorrelation stage.

5: The algorithm is stopped when $N_{\rm sim}$ reaches some user-chosen time $N_{\rm end}$. The stationary average $\pi(f)$ is estimated as

$$\pi(f) \approx F(f)_{\rm sim}/F(1)_{\rm sim}.$$

embedded ParRep is a Fleming-Viot based sampling procedure [3, 9]. The procedure can be summarized as follows. The R replicas $X_n^1, \ldots, X_n^R$, starting in W, evolve until one or more of them leaves W. Then each replica that left W is restarted from the current state of another replica that is currently in W (usually chosen uniformly at random). The procedure stops after the replicas have evolved for $n = n_p$ time steps, where n_p is a parameter similar to n_c. (If all the replicas leave W at the same time, the procedure restarts from the beginning.) With this type of sampling, the number of time steps simulated for each replica in the dephasing step is the same. In particular, if the R parallel processors are synchronous (i.e. if each processor takes the same wall clock time to simulate one time step), then each processor finishes the dephasing step at the same wall clock time.

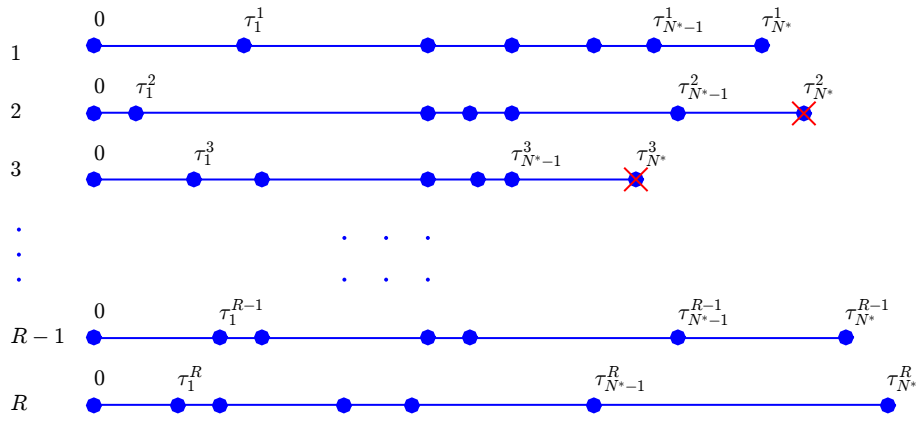

FIG. 2. *The diagram for embedded ParRep algorithm. Both replica* 2 *and replica* 3 *leave* W *after* $N^* = 6$ *transitions, in which case* $K = 2$.

- The acceleration of embedded ParRep comes from the parallel stage. Roughly, we only have to wait N^* time steps instead of N to observe an exit from W. The resulting wall clock time speedup can be approximately a factor of R. See Theorem 9 below.
  Like with CTMC ParRep, the parallel step does not require metastability for this time speedup, but if W is not metastable, then the dephasing step will not be efficient. See the remarks below Algorithm 1.
- To see that each parallel stage has a consistent contribution to the stationary average, we make the following two observations. Suppose that $x_1, \ldots, x_R$ are iid samples from ν. Then:
  1. The joint law of $(X^K_{N^K}, R(N^K-1)+K)$ is the same as that of $(X^1_{N^1}, N^1)$. That is, the joint distribution of the first exit time and the exit state for each parallel stage is independent of the number of replicas.
  2. The expected value of
  $$\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X^r_{N^*-1})\Delta\tau^r_{N^*-1}$$
  is the same as that of
  $$\sum_{n=0}^{N-1} f(X_n^1)\Delta\tau_n^1.$$
  That is, the expected contribution to $F(f)_{\text{sim}}$ from each parallel stage is independent of the number of replicas. See Theorem 11 below.

See Theorems 10 and 11 for proofs of these statements. Consistency of stationary averages follows from 1-2 above and the law of large numbers: since there are indefinitely many parallel stages in a given W, consistency is ensured so long as the expected contribution to $F(f)_{sim}$ from the parallel stage has the correct expected value. It isn't necessary that the contribution to $F(f)_{sim}$ from the parallel stage has the same *law* for any number of replicas – in fact, this turns out to be false for embedded ParRep.

We expect that embedded ParRep is superior to the CTMC ParRep for the following two reasons. First, consider the parallel stages of both algorithms. In CTMC ParRep, observing the first exit event in the parallel stage is not enough to determine T^*. But in embedded ParRep, once any replica leaves W, we know N^*. That is, the embedded ParRep parallel step terminates once any of the replicas leaves W. For this reason we expect the parallel stage of embedded ParRep to be significantly faster than that of CTMC ParRep.

Second, consider the dephasing stage. For embedded ParRep, Fleming-Viot sampling is a natural technique because, if the processors are synchronous, then they all finish the dephasing stage at the same wall clock time, and only the current states of each processor are needed at each time step to decide where to restart replicas which left W. (For asynchronous processors, one can simply implement a polling time.) This is not true, however, for Fleming-Viot sampling with CTMC ParRep. Indeed, to implement Fleming-Viot sampling with CTMC ParRep, one would have to store the histories of every replica, and the replicas would finish at potentially very different wall clock times.

The rejection method can be slow for both algorithms, particularly when the metastability is weak or when the number of replicas is large.

4.2. Error Analysis of Embedded ParRep. Now we are able to show that if the dephasing sampling is exact, then the expected contribution to $F(f)_{\text{sim}}$ from the parallel stage is exact.

THEOREM 11. *Suppose in the dephasing step $(x_1, \dots, x_R) \sim \mu^R$. Then the expected contribution to $F(f)_{sim}$ from the parallel stage of Algorithm 2 is the same for every number of replicas:*

$$\mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r\right] = \mathbb{E}^{\mu}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] = \mu(fq^{-1})\mathbb{E}^{\mu}\left[N\right].$$

Proof. We first rewrite

$$\begin{aligned}&\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r \\ &= \sum_{r=1}^{R}\sum_{i=0}^{N^*-1} f(X_i^r)\Delta\tau_i^r - \sum_{r=K+1}^{R} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r.\end{aligned} \tag{16}$$

For the first part, we condition N^* and obtain

$$\mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{i=0}^{N^*-1} f(X_i^r)\Delta\tau_i^r\right] = \sum_{r=1}^{R}\sum_{n=1}^{\infty}\sum_{i=0}^{n-1}\mathbb{E}^{\mu^R}\left[f(X_i^r)\Delta\tau_i^r\mathbb{I}_{N^*=n}\right]$$

Interchanging the iterated summations leads to

$$\sum_{r=1}^{R}\sum_{n=1}^{\infty}\sum_{i=0}^{n-1}\mathbb{E}^{\mu^R}\left[f(X_i^r)\Delta\tau_i^r\mathbb{I}_{N^*=n}\right] = \sum_{r=1}^{R}\sum_{i=0}^{\infty}\mathbb{E}^{\mu^R}\left[f(X_i^r)\mathbb{I}_{N^*>i}\Delta\tau_i^r\right].$$

Notice $N^* > i$ is equivalent to $N^1 > i, \dots, N^R > i$ and $\Delta\tau_i^r$ is independent of N^s for

$s \neq r$. Thus

$$\begin{aligned}&\sum_{r=1}^{R}\sum_{i=0}^{\infty}\mathbb{E}^{\mu^R}\left[f(X_i^r)\Delta\tau_i^r|N^*>i\right]\mathbb{P}^{\mu^R}(N^*>i)\\ =&\sum_{r=1}^{R}\sum_{i=0}^{\infty}\mathbb{E}^{\mu^R}\left[f(X_i^r)\Delta\tau_i^r|N^r>i\right]\mathbb{P}^{\mu^R}(N^*>i).\end{aligned}$$

Now by Lemma 1 and definition of the QSD,

$$\begin{aligned}\mathbb{E}^{\mu}\left[f(X_i^r)\Delta\tau_i^r|N^r>i\right]&=\mathbb{E}^{\mu}\left[\mathbb{E}^{\mu}\left[f(X_i^r)\Delta\tau_i^r\,|\,\{X_n^r\}_{n=0,1,\ldots}\right]|\,N^r>i\right]\\&=\mathbb{E}^{\mu}\left[f(X_i^r)\mathbb{E}^{\mu}\left[\Delta\tau_i^r\,|\,\{X_n^r\}_{n=0,1,\ldots}\right]|\,N^r>i\right]\\&=\mathbb{E}^{\mu}\left[f(X_i^r)q(X_i^r)^{-1}\big|\,N^r>i\right]=\mu(fq^{-1}).\end{aligned}$$

Combining the last four equations gives

$$\mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{i=0}^{N^*-1}f(X_i^r)\Delta\tau_i^r\right]=\mu(fq^{-1})R\mathbb{E}^{\mu^R}[N^*]. \tag{17}$$

A similar argument can be applied to the second term on the right hand side of (16). First we condition N^* and K simultaneously such that

$$\begin{aligned}&\mathbb{E}^{\mu^R}\left[\sum_{r=K+1}^{R}f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r\right]\\=&\sum_{n=1}^{\infty}\sum_{r=1}^{R}\sum_{r=k+1}^{R}\mathbb{E}^{\mu^R}\left[f(X_{n-1}^r)\Delta\tau_{n-1}^r|N^*=n,K=k\right]\mathbb{P}^{\mu^R}(N^*=n,K=k).\end{aligned}$$

Interchanging the second and third summations, the right-hand side equals

$$\sum_{n=1}^{\infty}\sum_{r=2}^{R}\sum_{k=1}^{r-1}\mathbb{E}^{\mu^R}\left[f(X_{n-1}^r)\Delta\tau_{n-1}^r|N^*=n,K=k\right]\mathbb{P}^{\mu^R}(N^*=n,K=k)$$

Recall that

$$N^*=n,\,K=k\Longleftrightarrow N^1>n,\ldots,N^{k-1}>n,\,N^k=n,\,N^{k+1}>n-1,\ldots,N^R>n-1.$$

Thus, using independence of $X_n^1,\ldots,X_n^R$ and definition of QSD,

$$\begin{aligned}&\sum_{n=1}^{\infty}\sum_{r=2}^{R}\sum_{k=1}^{r-1}\mathbb{E}^{\mu^R}\left[f(X_{n-1}^r)\Delta\tau_{n-1}^r|N^*=n,K=k\right]\mathbb{P}^{\mu^R}(N^*=n,K=k)\\=&\sum_{n=1}^{\infty}\sum_{r=2}^{R}\sum_{k=1}^{r-1}\mathbb{E}^{\mu}\left[f(X_{n-1}^r)\Delta\tau_{n-1}^r|N^r>n-1\right]\mathbb{P}^{\mu^R}(N^*=n,K=k)\\=&\mu(fq^{-1})\sum_{n=1}^{\infty}\sum_{r=2}^{R}\sum_{k=1}^{r-1}\mathbb{P}^{\mu^R}(N^*=n,K=k)\\=&\mu(fq^{-1})(R-\mathbb{E}^{\mu^R}[K]).\end{aligned}$$

Combining the last three equations leads to

$$\mathbb{E}^{\mu}\left[\sum_{r=K+1}^{R} f(X^r_{N^*-1})\Delta\tau^r_{N^*-1}\right] = \mu(fq^{-1})(R - \mathbb{E}^{\mu^R}[K]). \tag{18}$$

Subtracting (18) from (17), we have

$$\mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{i=0}^{N^*-1} f(X^r_i)\Delta\tau^r_i - \sum_{r=K+1}^{R} f(X^r_{N^*-1})\Delta\tau^r_{N^*-1}\right] = \mu(fq^{-1})\mathbb{E}^{\mu^R}\left[R(N^*-1)+K\right].$$

Now the result follows since

$$\mu(fq^{-1})\mathbb{E}^{\mu^R}\left[R(N^*-1)+K\right] = \mu(fq^{-1})\mathbb{E}^{\mu}[N]$$

by the Theorem 10. In particular, when $R = 1$ we have $N^* = N$ and $K = 1$, and thus

$$\mathbb{E}^{\mu}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] = \mu(fq^{-1})\mathbb{E}^{\mu}[N].$$

∎

We now prove an analog of Theorem 8 for embedded ParRep. Recall we have assumed convergence of $\|\mu_{n_c} - \mu\|_{TV} \to 0$ as $n_c \to \infty$, for every starting point $x \in E$, where $\|\cdot\|$ denotes total variation norm. See for instance Theorem 3 for conditions guaranteeing this convergence.

THEOREM 12. *Consider embedded ParRep starting at $X^{par}_{N_{sim}} = x \in W$ in the decorrelation stage, where WLOG $N_{sim} = 0$.*

Define

$$N^{par} = \begin{cases} N & \text{if } n_c > N \\ n_c + R(N^*-1) + K & \text{if } n_c \le N. \end{cases} \tag{19}$$

Assume the dephasing stage sampling is exact, that is, $(x_1, \dots, x_R) \sim \mu^R$. Consider the expected contribution to $F(f)_{sim}$ up until the first time the simulation leaves W (either in the decorrelation stage or in the parallel stage):

$$\Delta F(f)_{sim} \triangleq \mathbb{E}^{x,\mu^R}\left[\sum_{n=N_{sim}}^{N_{corr}-1} f(X^{par}_n)\Delta\tau^{par}_n + \mathbb{1}_{N^{par}>n_c}\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X^r_k)\Delta\tau^r_k + \sum_{r=1}^{K} f(X^r_{N^*-1})\Delta\tau^r_{N^*-1}\right],$$

where $\mathbb{E}^{x,\mu^R}$ denotes expectation for X^{par}_n and $X^1_n, \dots, X^R_n, \Delta\tau^1_n, \dots, \Delta\tau^R_n$ when $X^{par}_0 = x$ and $(X^1_n, \dots, X^R_n)$ has initial distribution μ^R. The error compared to direct (serial) simulation satisfies the bound

$$\left|\mathbb{E}^{x}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] - \Delta F(f)_{sim}\right| \le \|f\|_{\infty}\sup_{x\in W}\mathbb{E}^{x}[T]\,\|\mu_{n_c} - \mu\|_{TV}. \tag{20}$$

Proof of Theorem 12. Since X^{par}_n has the same law as X_n,

$$\mathbb{E}^{x}\left[\sum_{n=0}^{n_c\wedge N-1} f(X_n)\Delta\tau_n\right] = \mathbb{E}^{x,\mu^R}\left[\sum_{n=0}^{n_c\wedge N^{\mathrm{par}}-1} f(X^{\mathrm{par}}_n)\Delta\tau^{par}_n\right].$$

Thus,

$$
\begin{aligned}
&\left|\mathbb{E}^x\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] - \Delta F(f)_{sim}\right| \\
&= \left|\mathbb{E}^x\left[\sum_{n=n_c\wedge N}^{N-1} f(X_n)\Delta\tau_n\right]\right. \\
&\quad \left. - \mathbb{E}^{x,\mu^R}\left[\mathbb{1}_{N^{par}>n_c}\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r\right]\right| \\
&= \left|\mathbb{E}^x\left[\sum_{n=n_c}^{N-1} f(X_n)\Delta\tau_n \middle| N > n_c\right]\right. \\
&\quad \left. - \mathbb{E}^{x,\mu^R}\left[\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r \middle| N^{par} > n_c\right]\right| \mathbb{P}^\mu(N > n_c) \\
&\leq \left|\mathbb{E}^x\left[\sum_{n=n_c}^{N-1} f(X_n)\Delta\tau_n \middle| N > n_c\right]\right. \\
&\quad \left. - \mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r\right]\right|.
\end{aligned}
$$

By the Markov property,

$$
\mathbb{E}^x\left[\sum_{n=n_c}^{N-1} f(X_n)\Delta\tau_n \middle| N > n_c\right] = \mathbb{E}^{\mu_{n_c}}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right].
$$

By Theorem 11,

$$
\mathbb{E}^{\mu^R}\left[\sum_{r=1}^{R}\sum_{k=0}^{N^*-2} f(X_k^r)\Delta\tau_k^r + \sum_{r=1}^{K} f(X_{N^*-1}^r)\Delta\tau_{N^*-1}^r\right] = \mathbb{E}^{\mu}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right].
$$

Combining the last three equations,

$$
\begin{aligned}
&\left|\mathbb{E}^x\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] - \Delta F(f)_{sim}\right| \\
&\leq \left|\mathbb{E}^{\mu_{n_c}}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right] - \mathbb{E}^{\mu}\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right]\right| \\
&= \left|\sum_{x\in W}\mathbb{E}^x\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right]\mu_{n_c}(x) - \sum_{x\in W}\mathbb{E}^x\left[\sum_{n=0}^{N-1} f(X_n)\Delta\tau_n\right]\mu(x)\right| \\
&\leq \|f\|_\infty \sup_{x\in W}\mathbb{E}^x[T]\|\mu_{n_c} - \mu\|_{TV}
\end{aligned}
$$

with the last equation coming from the fact that $\mathbb{E}^x[\sum_{n=0}^{N-1}\Delta\tau_n] = \mathbb{E}^x[T]$. ∎

5. Numerical Experiment. We present an example from the stochastic reaction networks to demonstrate the consistency of the ParRep algorithms we proposed.

We consider the following stochastic reaction network

$$\varnothing \longrightarrow A \underset{}{\overset{1/\epsilon}{\rightleftharpoons}} B \longrightarrow C \longrightarrow \varnothing \tag{21}$$

from [7], where A, B and C represent the type of species and the scale disparity parameter $\epsilon \ll 1$. The time evolution of the population in the reaction network is commonly modeled by a CTMC $X(t) = (X_1(t), X_2(t), X_3(t))$ with state space $E \subset \mathbb{Z}_+^3$. By convention, we use the vector x_0 to denote the initial population. The jump rate of each reaction is governed by the intensity function $\lambda_j(x), j = 1, \ldots, 5$ such that for all $t > 0$,

$$\lambda_j(x) = \lim_{h \to 0} \frac{\mathbb{P}^{x_0}(X(t+h) = x + \eta_j | X(t) = x)}{h},$$

where η_j is the state change vector associated with the jth reaction. We list the reactions and their corresponding intensities and state change vectors in Table 1.

TABLE 1
Reactions, intensities and state change vectors

Reaction	Intensity	State change vector
$\varnothing \longrightarrow A$	$\lambda_1(x) = c_1$	$\eta_1 = (1, 0, 0)$
$A \longrightarrow B$	$\lambda_2(x) = \epsilon^{-1} c_2 x_1$	$\eta_2 = (-1, 1, 0)$
$B \longrightarrow A$	$\lambda_3(x) = \epsilon^{-1} c_3 x_2$	$\eta_3 = (1, -1, 0)$
$B \longrightarrow C$	$\lambda_4(x) = c_4 x_2$	$\eta_4 = (0, -1, 1)$
$C \longrightarrow \varnothing$	$\lambda_5(x) = c_5 x_3$	$\eta_5 = (0, 0, -1)$

Due to the scale disparity parameter ϵ, the reactions $A \to B$ and $B \to A$ occur with a much higher probability than the other reactions and hence we call $A \to B$ and $B \to A$ *fast reactions* and the other reactions *slow reactions*. The occurrence of slow reactions is a rare event. If we define the observables $f_1(x) = x_1 + x_2$ and $f_2(x) = x_3$, then the collection of sets $\{W_{m,n}\}_{m,n \in \mathbb{Z}_+}$ with

$$W_{m,n} = \{x \in E : f_1(x) = m, f_2(x) = n\}$$

form a full decomposition of the state space E. Note that both the total population of species A and B (i.e., $f_1(X(t))$) and the population of species C (i.e. $f_2(X(t))$) remain constants until one of the slow reactions occurs. Hence the typical sojourn time for $X(t)$ in each $W_{m,n}$ is very long comparing to the transition time between any two states that are in $W_{m,n}$, in which case, we say $X(t)$ is metastable in $W_{m,n}$. For example, with the initial population $x_0 = (1, 1, 0)$, the states $(1, 1, 0), (2, 0, 0)$ and $(0, 2, 0)$ form a metastable set since the fast reactions $A \to B$ and $B \to A$ occur with a significantly higher probability than slow reactions do and only the occurrence of the slow reactions can allow the process to move from the metastable set to another metastable set. Note that both observables f_1 and f_2 defined above are invariant in each metastable set, we call them slow observables. In general, an observable f is called a *slow observable* if it is invariant in each matestable set $W_{m,n}$, i.e., there is a constant $C(m, n)$ such that $f(x) = C(m, n)$ for each $x \in W_{m,n}$. An observable is called a *fast observable* if it is not slow (e.g., $f(x) = x_1$).

This kind two-scale problems arise in many fields other than the stochastic reaction networks, such as the queuing theory and population dynamics. Estimation of

the distributions of this kind two-scale processes can be computationally prohibitive due to the insufficient sampling of rare events. In the non-equilibrium regime, the distribution can be approximated using the stochastic averaging technique [19, 11] or tau-leaping scheme [14]. However, the former fails to apply in the equilibrium regime and the latter can be still computationally expensive for long time simulation In this situation, it is desirable to apply the two ParRep algorithms proposed in this paper to accelerate the long time simulation and estimate the stationary distribution.

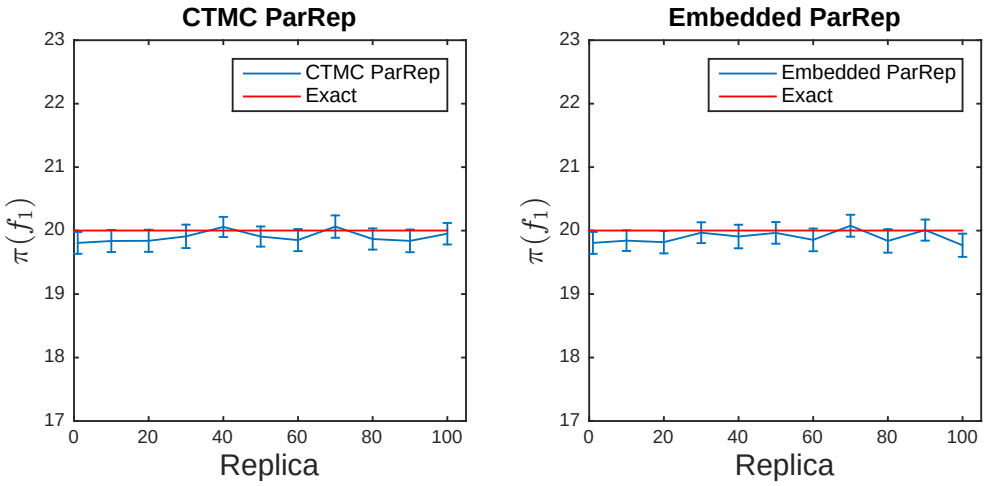

FIG. 3. *Stationary average of the slow observable f_1 computed with the CTMC ParRep (left) and with the embedded ParRep (right). The user-specified terminal time is $T_{end} = 10^4$ in the simulation.*

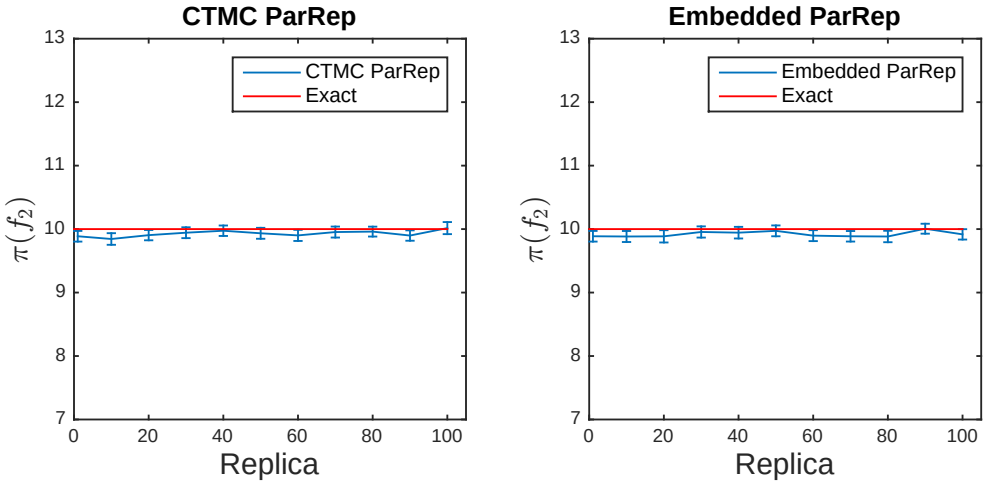

FIG. 4. *Stationary average of the slow observable f_2 computed with the CTMC ParRep (left) and with the embedded ParRep (right). The user-specified terminal time is $T_{end} = 10^4$ in the simulation.*

In this numerical experiment, we take the initial state $x_0 = (5, 10, 10)$ and the

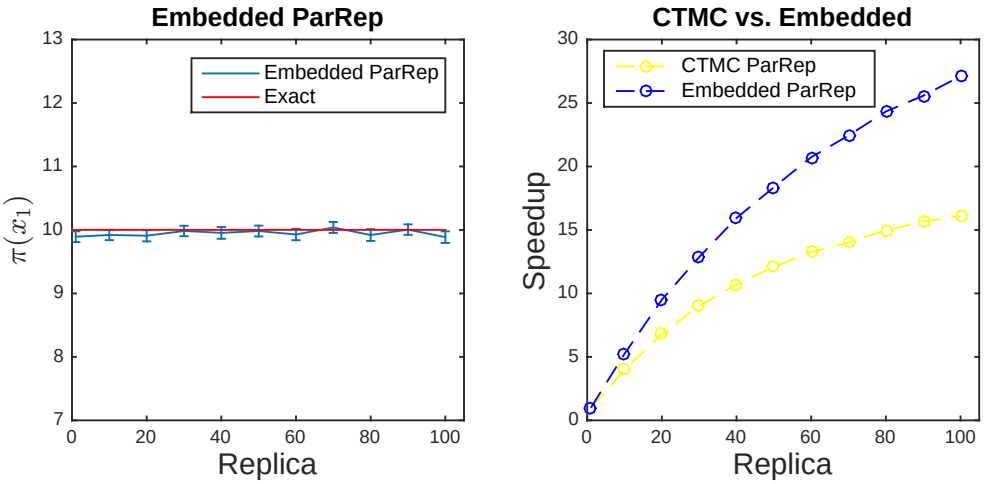

FIG. 5. *Stationary average of the fast observable x_1 computed with embedded ParRep (left) and the speedup comparison between the CTMC ParRep and the embedded ParRep (right). The user-specified terminal time is $T_{end} = 10^4$ in the simulation.*

rate parameters

$$(c_1, c_2, c_3, c_4, c_5) = (0.1, 100, 100, 0.01, 0.01).$$

With this choice of parameters, the scale disparity of this CTMC is of the order 10^4 and hence the process $X(t)$ demonstrates metastability. In this case, simulation of this network is computationally demanding since one needs to run the SSA for very long time for the process to be well into the equilibrium state in order to estimate the stationary distribution. We apply both the CTMC ParRep and the embedded ParRep to estimate the stationary averages of the slow observables f_1 and f_2. The stationary distribution of the fast observable $f(x) = x_1$ is also computed using the embedded ParRep. On the other hand, for the reaction network (21) under consideration, one can calculate the stationary distribution analytically since it only involves mono-molecular reactions. In fact, it can be shown that the stationary distribution is a multivariate Poisson distribution [7], that is,

$$\pi(x_1, x_2, x_3) = \frac{\bar{\lambda}_1^{x_1} \bar{\lambda}_2^{x_2} \bar{\lambda}_3^{x_3}}{x_1! x_2! x_3!} e^{-(\bar{\lambda}_1 + \bar{\lambda}_2 + \bar{\lambda}_3)}, \tag{22}$$

where

$$\bar{\lambda}_1 = \frac{c_1(c_3 + c_4)}{c_2 c_4}, \quad \bar{\lambda}_2 = \frac{c_1}{c_4}, \quad \bar{\lambda}_3 = \frac{c_1}{c_5}.$$

Hence the exact stationary averages of the slow observables f_1 and f_2 are $\pi(f_1) = 20.001$ and $\pi(f_2) = 10$ and the exact stationary averages of the fast observable $f(x) = x_1$ is 10.001. We use this exact result to compare with our result from numerical simulation.

Our simulations compare the CTMC ParRep and the embedded ParRep with the SSA. In Figure 3, we demonstrate the estimation of $\pi(f_1)$ using CTMC ParRep and embedded ParRep with various number of replicas ($R = 10, 20, \cdots, 100$) and with SSA ($R = 1$). Similarly, Figure 4 shows the estimation of $\pi(f_2)$. Note that only the embedded ParRep is used to compute the stationary average of the fast

variable $f(x) = x_1$ since the CTMC ParRep is not efficient for fast observables as we commented at the end of the Section 3.1. Currently, the rejection sampling is used for dephasing and the decorrelation and dephasing thresholds are taken to be $t_c = t_p = 0.01$ for CTMC ParRep and $n_c = n_p = 15$ steps for embedded ParRep. In Figure 5, the estimation for the fast observable and speedup are shown. It can be seen that with 10 replicas, the speedup factor is about 4.5 for the CTMC ParRep and 5.5 for the embedded ParRep. When the number of replicas increases, the embedded ParRep becomes much more efficient than the CTMC ParRep. However, even the embedded ParRep is far away from the linear speedup (with 100 replicas, about 27 times faster than SSA). This sublinear speedup comes from the fact that when the number of replica is large, the acceleration is offset by the inefficient rejection sampling based dephasing procedure. We expect that the embedded ParRep would be more efficient if the Fleming-Viot particle processes are used for dephasing.

Acknowledgment. The work of T.W. has been supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under the contract number DE-SC0010549. The work of P.P. was partially supported by the DARPA project W911NF-15-2-0122. D.A. gratefully acknowledges support from the National Science Foundation via the award NSF-DMS-1522398.

REFERENCES

[1] D. Aristoff, *The parallel replica method for computing equilibrium averages of Markov chains*, Monte Carlo Methods and Applications, 21 (2015), pp. 255–273.

[2] D. Aristoff, T. Lelièvre, and G. Simpson, *The parallel replica method for simulating long trajectories of markov chains*, Applied Mathematics Research eXpress, 2014 (2014), pp. 332–352.

[3] A. Binder, T. Lelièvre, and G. Simpson, *A generalized parallel replica dynamics*, Journal of Computational Physics, 284 (2015), pp. 595–616.

[4] J. Blanchet, P. Glynn, and S. Zheng, *Theoretical analysis of a stochastic approximation approach for computing quasi-stationary distributions*, arXiv preprint arXiv:1401.0364, (2014).

[5] P. Brémaud, *Markov Chains: Gibbs fields, Monte Carlo simulation, and Queues*, Springer, 1998.

[6] P. Collet, S. Martínez, and J. San Martín, *Quasi-stationary distributions: Markov chains, diffusions and dynamical systems*, Springer Science & Business Media, 2012.

[7] S. L. Cotter, K. C. Zygalakis, I. G. Kevrekidis, and R. Erban, *A constrained approach to multiscale stochastic simulation of chemically reacting systems*, The Journal of chemical physics, 135 (2011), p. 094102.

[8] P. Del Moral and A. Doucet, *Particle motions in absorbing medium with hard and soft obstacles*, Stochastic Analysis and Applications, 22 (2004), pp. 1175–1207.

[9] P. A. Ferrari and N. Maric, *Quasi stationary distributions and fleming-viot processes in countable spaces*, Electron. J. Probab, 12 (2007), pp. 684–702.

[10] D. T. Gillespie, *Exact stochastic simulation of coupled chemical reactions*, The journal of physical chemistry, 81 (1977), pp. 2340–2361.

[11] A. Hashemi, M. Nunez, P. Plecháč, and D. G. Vlachos, *Stochastic averaging and sensitivity analysis for two scale reaction networks*, The Journal of chemical physics, 144 (2016), p. 074104.

[12] C. Le Bris, T. Lelievre, M. Luskin, and D. Perez, *A mathematical formalization of the parallel replica dynamics*, Monte Carlo Methods and Applications, 18 (2012), pp. 119–146.

[13] D. Perez, B. P. Uberuaga, and A. F. Voter, *The parallel replica dynamics method–coming of age*, Computational Materials Science, 100 (2015), pp. 90–103.

[14] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, *Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method*, The Journal of Chemical Physics, 119 (2003), pp. 12784–12794.

[15] R. Y. Rubinstein and D. P. Kroese, *Simulation and the Monte Carlo method*, vol. 707, John

Wiley & Sons, 2011.
[16] C. Schütte and M. Sarich, *Metastability and Markov State Models in Molecular Dynamics*, Courant Lecture Notes, 2013.
[17] E. Seneta, *Non-negative matrices and Markov chains*, Springer Science & Business Media, 2006.
[18] A. F. Voter, *Parallel replica method for dynamics of infrequent events*, Physical Review B, 57 (1998), p. R13985.
[19] E. Weinan, D. Liu, and E. Vanden-Eijnden, *Nested stochastic simulation algorithms for chemical kinetic systems with multiple time scales*, Journal of computational physics, 221 (2007), pp. 158–180.